

BLAKELY SOKOLOFF TAYLOR & ZAFMAN LLP

TELEPHONE: (408) 720-8300

INTELLECTUAL PROPERTY LAW
12400 WILSHIRE BOULEVARD, 7TH FLOOR
LOS ANGELES, CA 90025

FACSIMILE: (408) 720-8383

RECEIVED
CENTRAL FAX CENTER**MAY 03 2006****FACSIMILE COVER SHEET**

Deliver to: Shin, Christopher B., USPTO Art Group: 2182
 Facsimile No.: (571) 273-8300 Date: May 3, 2006
 From: Chui-Kiu Teresa Wong, Reg. No. 48,042
 Our Docket No.: 42390P6344 Number of pages 44 including this sheet.
 Application No.: 09/186,056 Filing Date: 11/3/1998
 Docket Due Date(s): _____

Enclosed are the following documents:

- | | |
|--|--|
| <input type="checkbox"/> Amendment: _____ (____ pgs) | <input type="checkbox"/> Issue Fee Transmittal |
| <input type="checkbox"/> Appeal Brief (____ pgs). | <input type="checkbox"/> Notice of Appeal |
| <input type="checkbox"/> Application: _____ | <input checked="" type="checkbox"/> Petition for: <u>Correction of Filing Date</u> |
| (____ pgs) w/cover & abstract | <input type="checkbox"/> Request for Continued Examination (RCE) |
| <input type="checkbox"/> Assignment & Cover Sheet (____ pgs) | <input type="checkbox"/> Reply Brief (____ pgs) |
| <input type="checkbox"/> Certificate of _____ | <input type="checkbox"/> Request & Certification Under 35 USC 122(b)(2)(B)(i) |
| <input type="checkbox"/> Continued Prosecution Application (CPA) | <input type="checkbox"/> Request to Rescind Previous Nonpublication Request |
| <input type="checkbox"/> Declaration & POA (____ pgs) | <input type="checkbox"/> Response to Notice of Missing Parts & Formalities Letter |
| <input type="checkbox"/> Drawings: ____ sheets, ____ figures | <input type="checkbox"/> Response to Written Opinion (____ pgs) |
| <input type="checkbox"/> Extension of Time: _____ | <input type="checkbox"/> Terminal Disclaimer |
| <input type="checkbox"/> Fee Transmittal (in duplicate) | <input type="checkbox"/> Transmittal of Publication Fee Due |
| <input type="checkbox"/> IDS & PTO/SB/08 (____ pgs) | <input type="checkbox"/> Transmittal Letter |
| <input checked="" type="checkbox"/> Other <u>Originally filed application, copy of postcard returned by the USPTO, previous petition filed and notice of withdrawal of previously sent notice.</u> | |

CERTIFICATE OF MAILING/TRANSMISSION (37 CFR 1.8A)

I hereby certify that this correspondence is being transmitted by facsimile on the date shown below to the United States Patent and Trademark Office.

Vanessa Sanchez
 Vanessa Sanchez

5/3/2006

Date

Confidentiality Note: The documents accompanying this facsimile transmission contain information from the law firm of Blakely, Sokoloff, Taylor & Zafman which is confidential or privileged. The information is intended to be for the use of the individual or entity named on this transmission sheet. If you are not the intended recipient, be aware that any disclosure, copying, distribution or use of the contents of this faxed information is prohibited. If you have received this facsimile in error, please notify us by telephone immediately so that we can arrange for the retrieval of the original documents at no cost to you.

If you do not receive all the pages, or if there is any difficulty in receiving, please call: (408) 720-8300 and ask for Vanessa Sanchez.

042390.P6344
09/186,056

Patent

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

RECEIVED
CENTRAL FAX CENTER

MAY 03 2006

In re Application of:

David I. Poisner, et al.

Application No.: 09/186,056

Filed: November 3, 1998

Attorney Docket No.: 042390.P6344

For: RACE FREE DATA TRANSFER
ALGORITHM USING HARDWARE
BASED POLLING

Examiner: Shin, Christopher B.

Art Unit: 2182

Office of Petitions
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

PETITION TO CORRECT FILING DATE
37 C.F.R. § 1.181(a)

In response to the Notice of Allowance mailed February 9, 2006, Applicants respectfully request the correction of the filing date in view of the remarks herein:

CERTIFICATE OF TRANSMISSION VIA FACSIMILE

I hereby certify that this correspondence is being facsimile transmitted to the U.S. Patent and Trademark Office on: May 3, 2006.

Vanessa Sanchez
(Typed or printed name of person transmitting paper)


(Signature of person transmitting paper)

042390.P6344
09/186,056

Patent

STATEMENTS OF FACTS

Applicant hereby petitions to request correction of the filing date of the current application. The current application was originally filed on November 3, 1998 with a specification, drawings, a list of claims, and the appropriate fee. Since Applicant never received the notice of missing parts, and therefore, was unable to respond, the application was held to be abandoned and a notice of abandonment was mailed on October 28, 2003. Accordingly, Applicant timely filed a petition on December 19, 2003, which the USPTO received on December 22, 2003. The petition was granted and a notice of withdrawal of previously sent notice was mailed on December 1, 2004. However, instead of according the original filing date (November 3, 1998) to the current application, the date the USPTO received the previously petition (December 22, 2003) was erroneously designated as the filing date of the current application.

042390.P6344
09/186,056

Patent

ACTION REQUESTED

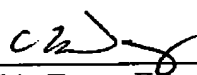
Applicant therefore requests that the filing date of the current application be corrected to be November 3, 1998.

Applicant has submitted herewith a copy of the originally filed application, the postcard returned by the USPTO, the previous petition filed, and the notice of withdrawal of previously sent notice.


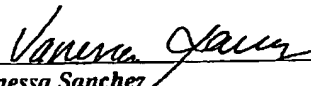
If there is a deficiency in fees, please charge our Deposit Account No. 02-2666.

Respectfully submitted,
BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP

Date: 5/3, 2006


Chui-kiu Teresa Wong
Attorney for Applicant
Reg. No. 48,042

12400 Wilshire Boulevard
Seventh Floor
Los Angeles, California 90025-1026
(408) 720-8300

REQUEST FOR CORRECTED DATE OF FILING		Docket Number 42390P6344						
In re Application of David Poisner, et al.								
Application Number 09/186,056	Filed 11/3/1998							
For RACE FREE DATA TRANSFER ALGORITHM USING HARDWARE BASED POLLING								
Group Art Unit 2182	Examiner Shin, Christopher B.							
<p>1. Attached is a copy of the Originally filed application, returned postcard, previously filed petition, notice of withdrawal of previously sent notice received from the PTO in the above application for which issuance of a corrected Originally filed application, returned postcard, previously filed petition, notice of withdrawal of previously sent notice is respectfully requested.</p> <p>2. There is an error with respect to the following data, which is:</p> <p><input checked="" type="checkbox"/> Incorrectly entered</p> <p>and/or</p> <p><input type="checkbox"/> omitted.</p> <table border="1"><thead><tr><th>Error in</th><th>Correct Data</th></tr></thead><tbody><tr><td>Filing Date</td><td>November 3, 1998</td></tr><tr><td>December 22, 2003</td><td></td></tr></tbody></table>			Error in	Correct Data	Filing Date	November 3, 1998	December 22, 2003	
Error in	Correct Data							
Filing Date	November 3, 1998							
December 22, 2003								
<p>Respectfully submitted,</p> <p>BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP</p> <p>Dated: <u>May 3, 2006</u></p> <p> Chui-Kiu Teresa Wong, Reg. No. 48,042</p> <p>12400 Wilshire Boulevard, 7th Floor Los Angeles, CA 90025 Telephone: (408) 720-8300</p> <p>CERTIFICATE OF MAILING/TRANSMISSION I hereby certify that this correspondence is being transmitted via facsimile on the date shown below to the United States Patent and Trademark Office.</p> <p> Vanessa Sanchez</p> <p>05/03/06 Date</p>								

156



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING OR 371(c) DATE	FIRST NAMED APPLICANT	ATTY. DOCKET NO./TITLE
09/186,056		DAVID I. POISNER	042390.P6344


BLAKLEY SOKOLOFF TAYLOR & ZAFMAN
12400 WILSHIRE BOULEVARD
7TH FLOOR
LOS ANGELES, CA 90025

CONFIRMATION NO. 1348
WITHDRAWAL NOTICE
OC000000012864390
OC000000012864390

Date Mailed: 12/01/2004

WITHDRAWAL OF PREVIOUSLY SENT NOTICE

The Notice of Abandonment mailed on 10/28/2003 was sent in error and is hereby withdrawn. A Filing Receipt is enclosed. The Office regrets any inconvenience the error may have caused.



RECEIVED

DEC 06 2004

BLAKLEY, SOKOLOFF, TAYLOR & ZAFMAN LLP
LOS ANGELES

*A copy of this notice **MUST** be returned with the reply.*

Deborahano-Lyles
Customer Service Center
Initial Patent Examination Division (703) 308-1202

PART 1 - ATTORNEY/APPLICANT COPY

Docket No.: 42390.P6344

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re the application of:)	
)	
Poisner et al.)	
)	
Serial No.: 09/186,056)	Examiner: TBD
)	
Filed: 11/3/98)	Art Unit: TBD
)	
For: RACE FREE DATA TRANSFER)	
ALGORITHM USING HARDWARE)	
BASED POLLING)	

PETITION TO WITHDRAW HOLDING OF ABANDONMENT BASED ON FAILURE
TO RECEIVE OFFICE COMMUNICATION
37 CFR § 1.181(a)

Hon. Commissioner of
Patent & Trademarks
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Sir:

In response to the Notice of Abandonment mailed 10/28/03, applicants respectfully requests withdrawal of holding of abandonment for the above referenced application in view of applicant's remarks herein.

42390.P6344
Serial #: 09/186,056

-1-

STATEMENT OF FACTS

Applicant hereby petitions to request withdrawal of the holding of abandonment mailed 10/28/2003. The abandonment was held due to an alleged failure to respond to a Notice of Missing Parts. As detailed below, applicants never received this notice of missing parts and therefore were unable to respond. Applicant submits herewith the appropriate response to the Notice of Missing Parts as well as the proof required to justify withdrawal of the holding of abandonment specified in the MPEP.

According to MPEP § 711.03(c), the showing required to establish nonreceipt of an Office communication must include a statement from the practitioner stating that the Office communication was not received by the practitioner and attesting to the fact that a search of the file jacket and docket records indicates that the Office communication was not received. A copy of the docket record where the nonreceived Office communication would have been entered had it been received and docketed must be attached to and referenced in practitioner's statement.

Accordingly, it is hereby submitted that:

1. A search of the above-referenced patent file, Intel docket number P6344, Ser. No. 09/186,056 was performed, and indicated that the Notice of Missing Parts, allegedly mailed by the Patent Office on 11/19/98 was never received because it was not present in the file.
2. Notably, a request for status was sent to the Patent Office requesting information about the status of the case on 11/6/2000 due to the long latency without any communications being received from the Patent Office.
3. A search of the docket records pertinent to this file was also performed and also

42390.P6344
Serial #: 09/186,056

-2-

indicated that the Notice of Missing Parts was never received by applicants. A copy of the docket record where the non-received Notice of Missing Parts would have been entered had it been received and docketed is attached.

ACTION REQUESTED

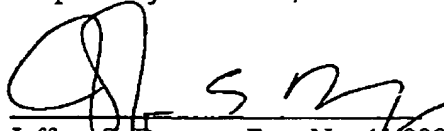
Applicants therefore request that the holding of abandonment based on applicants' alleged failure to respond to the Notice of Missing Parts be withdrawn. Applicant submits that this request was timely submitted within 2 months of the notice of abandonment mailed 10/28/03.

Applicants have submitted herewith the signatures required for this case, although since no notice of missing parts was received, no copy of such notice is enclosed.

If there is a deficiency in fees, please charge our Deposit Acct. No. 02-2666.

Date: 12/19/03

Respectfully submitted,


Jeffrey S. Draeger, Reg. No. 44,000

12400 Wilshire Boulevard
Seventh Floor
Los Angeles, CA 90025-1026
(408) 720-8598

42390.P6344
Serial #: 09/186,056

-3-

FIRST CLASS CERTIFICATE OF MAILING
(37 C.F.R. § 1.8(a))

I hereby certify that this correspondence is being deposited with the United States Postal Service as first class mail with sufficient postage in an envelope addressed to the Commissioner of Patents and Trademarks, Washington, D.C. 20231

on 12/19/03
Date of Deposit

Patricia M. Richard

Name of Person Mailing Correspondence

[Signature]
Signature

12/19/03
Date

42390.P6344
Serial #: 09/186,056

-4-

ENTERED

NOV 13 1998

STATUS DB-LA

RECEIVED
NOV 13 1998

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN
LOS ANGELES

Serial/Patent No.: **
Client: INTEL CORPORATION
Title: a race free data transfer algorithm using hardware based polling

BSTZ File No.: 042390.P6344
Date Mailed: 11/3/98
Filing/Issue Date: Herewith
Atty/Secty Initials: EHT/JSD/jkm
Docket Due Date: **

The following has been received in the U.S. Patent & Trademark Office on the date stamped hereon:

<input type="checkbox"/> Amendment/Response (____ pgs.)	<input type="checkbox"/> Express Mail No. EM335898998US	<input checked="" type="checkbox"/> Check No. 250-78
<input type="checkbox"/> Appeal Brief (____ pgs.) (in triplicate)	<input type="checkbox"/> Month(s) Extension of Time	Amt: 1020.00
<input type="checkbox"/> Application - Rule 1.53(b) Divisional (____ pgs.)	<input type="checkbox"/> Information Disclosure Statement & PTD 149 (____ pgs.)	<input type="checkbox"/> Check No. _____
<input type="checkbox"/> Application - Rule 1.53(b) CIP (____ pgs.)	<input type="checkbox"/> Issue Fee Transmittal	Amt: _____
<input type="checkbox"/> Application - Rule 1.53(d) CPA Transmittal (____ pgs.)	<input type="checkbox"/> Notice of Appeal	
<input type="checkbox"/> Application - Design (____ pgs.)	<input type="checkbox"/> Petition for Extension of Time	
<input type="checkbox"/> Application - FCT (____ pgs.)	<input type="checkbox"/> Petition for _____	
<input type="checkbox"/> Application - Provisional (____ pgs.)	<input type="checkbox"/> Postcard	
<input type="checkbox"/> Assignment and Cover Sheet	<input type="checkbox"/> Power of Attorney (____ pgs.)	
<input type="checkbox"/> Certificate of Mailing	<input type="checkbox"/> Preliminary Amendment (____ pgs.)	
<input checked="" type="checkbox"/> Declaration & POA (4 pgs.) unsigned	<input type="checkbox"/> Reply Brief (____ pgs.)	
<input type="checkbox"/> Disclosure Docs & Orig & Copy of Invention Signed Letter (____ pgs.)	<input type="checkbox"/> Response to Notice of Missing Parts	
<input checked="" type="checkbox"/> Drawings: 6 # of sheets includes 6 figures	<input type="checkbox"/> Small Entity Declaration for Indep. Inventor/Small Business	
<input type="checkbox"/> Other: _____	<input checked="" type="checkbox"/> Transmittal Letter, in duplicate	
	<input checked="" type="checkbox"/> Fee Transmittal, in duplicate	

526 U.S. PTO
09/186056
11/03/98

RECEIVED

NOV 16 1998

S.T.Z. DATABASE DEPT

to F/F Socketed

UNITED STATES UTILITY PATENT APPLICATION

FOR

A RACE FREE DATA TRANSFER ALGORITHM USING HARDWARE BASED
POLLING

Inventors:

David I. Poisner
Karthi R. Vadivehu

42390.P6344

Prepared by:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN

12400 Wilshire Boulevard, Seventh Floor
Los Angeles, California 90025-1026
(408) 720-8598

EXPRESS MAIL CERTIFICATE OF MAILING

"Express Mail" mailing label number EM935898998 US
Date of Deposit 11-3-98

I hereby certify that this paper or fee is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service on the date indicated above and is addressed to the Commissioner of Patents and Trademarks, Washington, D.C. 20231.

Annette Jacobs
(Typed or printed name of person mailing paper or fee)

Annette Jacobs
(Signature of person mailing paper or fee)

A RACE FREE DATA TRANSFER ALGORITHM USING HARDWARE BASED POLLING

5

BACKGROUND

1. Field of the Invention

The present invention pertains to the field of data transfers in a computer or other processing system.

10

2. Description of Related Art

An improved data transfer algorithm, such as an improved direct memory access (DMA) technique, may provide advantages both in terms of improved system performance and in terms of ease of interface software (e.g., device driver) design. Improved system performance may result from fewer bus transactions being used and from lengthening the total data transfer when additional data becomes available during the course of the data transfer. Simpler device drivers may be designed if the device driver software can easily obtain a precise and updated indication of the transfer status throughout the transfer.

A DMA transfer is a transfer of data (i.e., any stored information, instructions, etc.) between system memory and a device with limited or no intervention from the system processor once the transfer commences. A memory region that acts as the source or target of a DMA transfer is often physically contiguous. Alternatively, some DMA controllers may allow access to scattered memory regions (i.e., they support scatter-gather). In a

42390.P6344

-2-

DMA controller supporting scatter-gather, either multiple addresses may be programmed into the DMA controller or a data structure may be used to track the multiple regions of memory.

One advantage of using DMA-style transfers is that a large block of memory may be automatically transferred without further intervention of the processor. In other words, the controller can be initialized, and then can provide numerous bus cycles to transfer data without further intervention. The controller, however, typically only proceeds until reaching an endpoint programmed in during the initialization. Additionally, some controllers provide no mechanism to notify other components or software routines of progress throughout the transfer.

As a result, inefficient latencies may develop and the data transfer process may be prematurely halted. For example, consider a transfer of buffers from memory to a DMA device. If there is no notification until all data from memory is transferred to the DMA device, then the space in memory used by the data is not released until the entire transfer is complete. This delayed release inefficiently reserves memory despite the fact that its contents may no longer be needed after the data transfer is complete. Moreover, additional data may have been prepared and placed in memory during the DMA transfer. If the DMA controller was aware of this data, it could also be transferred without interruption and re-initialization of the DMA controller.

Some prior art DMA techniques, however, do allow updating a value indicating the last buffer to be transferred during the DMA transfer. A stop bit or count may be stored at some point in memory (e.g., a stop bit may be within the buffer structure). When additional buffers become available for transmission, the last stop bit may be updated by

the software routine transferring the data into the additional buffers.

One problem with using a memory based stop bit is that the software routine(s) adding buffers to the list may experience a race condition with the DMA controller. Such software routines typically do not know exactly which buffer the DMA controller is
5 working on at a particular point in time. Therefore, there is a risk that a memory stop bit will be turned off by software after the DMA controller has already read the buffer and retrieved the enabled stop bit.

To overcome this race condition, one prior art approach requires the DMA controller to poll the last buffer indicator in memory. Such continuous polling may
10 disadvantageously use a large number of unnecessary bus cycles to read the pointer from memory. Additionally, such polling may still produce undesirable latencies. For example, if an additional buffer becomes ready for a transfer just after the DMA controller polls the value, the DMA controller will act on stale information until the next poll. As a result, the DMA controller may terminate a data transfer unnecessarily or at least experience a latency
15 until the next polling event since it is unaware that the additional buffer is ready.

Thus, the prior art may not provide an adequate data transfer technique. Some prior art techniques may either not allow additional transfers to be added after controller initialization, may not allow efficient independent preparation and/or reclamation of buffers, or may inject undesirable extra latencies or bus cycles.

20

Summary

A method and apparatus for a race free data transfer algorithm using hardware based polling is disclosed. One disclosed method transfers information between a target
5 device and a buffer which is one of a set of buffers. The buffer is pointed to by a current
buffer value stored in a controller. The current buffer value is adjusted to point to a next
buffer if the current buffer value is different than a last buffer value. One of the set of
buffers is serviced utilizing either the current buffer value or the last buffer value from the
controller.

10

15

42390.P6344

-5-

Brief Description of the Figures

The present invention is illustrated by way of example and not limitation in the
5 figures of the accompanying drawings.

Figure 1 illustrates one embodiment of a system utilizing disclosed data transfer
techniques.

10 Figure 2 illustrates a flow diagram for operations of one embodiment of the
system of Figure 1.

Figure 3 illustrates one embodiment of a system utilizing a buffer descriptor table.

15 Figure 4 illustrates a flow diagram for one embodiment of the system of Figure 3.

Figure 5 illustrates a flow diagram for one embodiment of a buffer reclamation
routine.

20 Figure 6 illustrates a flow diagram for one embodiment of a buffer preparation
routine.

Detailed Description

The following description provides a race free data transfer algorithm using
5 hardware based polling. In the following description, numerous specific details such as
register names, data structure configurations, specific system arrangements, and logic
partitioning and integration choices are set forth in order to provide a more thorough
understanding of the present invention. It will be appreciated, however, by one skilled in
the art that the invention may be practiced without such specific details. In other
10 instances, control structures and gate level circuits have not been shown in detail in order
not to obscure the invention. Those of ordinary skill in the art, with the included
descriptions, will be able to implement the necessary logic circuits without undue
experimentation.

The presently disclosed techniques may be used to achieve an improved DMA
15 architecture or to facilitate other similar types of data transfer transactions. A controller
utilizing disclosed techniques may use a current buffer value and a last buffer value to
allow the data transfer length to be changed after the transfer has been initiated without
causing a race condition to occur. In some embodiments, the presently disclosed
techniques may reduce the number of overhead bus cycles required to maintain a flexible
20 length transfer and/or may allow the use of simpler device driver software for interfacing
with the controller. Additionally, a prefetch register may be used to facilitate more
efficient accesses when a data structure used in the data transfer is stored in main
memory.

42390.P6344

-7-

Figure 1 illustrates one embodiment of a system using hardware based polling to maintain a race free DMA process. In Figure 1, a processor 145, a memory 150, and a DMA device 105 are coupled to a bus 140. More elaborate system architectures may be employed; however, the presently disclosed techniques are not limited to any particular system or set of device arrangements since data transfers may be implemented in a wide variety of ways. In fact, any system supporting a data transfer controller that may store values (e.g., in registers) which point to memory locations and are accessible to software routines initiating, controlling, or otherwise servicing the data transfer may be used for some embodiments.

In the embodiment shown in Figure 1, the DMA device 105 includes a DMA controller 100 and a DMA target 115. These components may be separate components or may be integrated together and/or with other components. The DMA target 115 has a data signal port 117. The DMA target 115 may send and/or receive either digital or analog signals via the port 117. Thus, the DMA device 105 may be almost any device which exchanges data with memory utilizing DMA or other similar transfers. Some examples are audio coder-decoders, modems, network interfaces, or other communication or signal exchange interfaces.

The DMA controller includes a current buffer register 110 and a last buffer register 112. Either a register as a dedicated storage location or a register as a particular entry in a general purpose storage area may be used. The current buffer register 110 and the last buffer register 112 may be implemented in a variety of manners so long as the values contained therein indicate or identify the proper buffer locations when needed. These registers allow the DMA controller 100 to track which one of a set of buffers 160

stored in the memory 150 is presently being transferred as well as the buffer at which the DMA controller 100 should stop (the last buffer).

For example, the registers may simply contain pointers to buffers in memory if a linked list type structure is used for the buffers. Alternatively, the registers may store either direct or indexed pointers into an optional data structure 155 stored in memory. If the optional data structure 155 is used, the registers may point to buffer descriptors residing in a buffer descriptor table in the memory 150. In any case, the values stored in the current buffer register 110 and the last buffer register 112 adequately indicate or identify particular buffers to the DMA controller 100 and any software routines that utilize this information.

The memory 150 also contains software routines. A buffer preparation routine 170 prepares buffers for a DMA transfer. If DMA data is being received by the memory 150 from the DMA target 115, then the buffer preparation routine 170 may assure that an empty buffer is ready to receive the data. If data is being transferred from the memory 150 to the DMA target 115, then the buffer preparation routine 170 may fill the buffer with data to be transferred.

A buffer reclamation routine 165 recaptures one or more buffers after the DMA transfer(s) affecting the buffer(s) complete. If DMA data is being received by the memory 150 from the DMA target 115, then the reclamation routine may pass the data on to the software process that requested the data and then may free the buffer. If data is being transferred from the memory 150 to the DMA target 115, then the information stored in the buffer is typically no longer needed, and the reclamation routine can mark the buffer as free. If additional DMA transfers complete before the reclamation routine

165 exits, the reclamation routine may continue to reclaim buffers.

The flow diagram of Figure 2 further illustrates operation of one embodiment of the system in Figure 1. In block 205, the current buffer register 110 and last buffer register 112 are initialized. In block 210, data is then transferred to or from the buffer indicated by the current buffer register 110.

After the transfer, an interrupt bit associated with the transferred buffer is tested in step 215. If the interrupt bit is set, then buffers are reclaimed as shown in block 225 via the buffer reclamation routine 165 (Fig. 1). Any processed buffer up to but excluding the current buffer may be reclaimed. If the buffer was a receiving buffer, the data may be transferred out and the buffer freed. If the buffer was a source buffer, the buffer may simply be marked free if the data is no longer needed.

Notably, utilizing an interrupt associated with each buffer is optional. As indicated by the dotted line 226, the reclamation routine could be an independent process which is periodically activated or which continuously frees processed buffers. In any case, the reclamation routine may be able to quickly free used buffers without waiting for the end of the entire DMA transfer because it can access the current buffer register 112 in the DMA controller and therefore determine which buffers are still needed.

In block 220, since processing of the current buffer has completed, the current buffer register 110 is incremented. Once again, the process returns to transferring data to or from the current buffer as shown in block 210. Accordingly, a series of buffers may be transferred using DMA-style data transfer techniques and the associated processed buffers may be freed for other uses.

Additionally, however, the length of the transfer may be increased by the buffer

42390.P6344

preparation routine 170 acting as a separate process 230 while the transfer is in progress. As shown in block 235, the buffer preparation routine 170 prepares buffers for transfer. This may involve designating a free buffer to receive data or filling a buffer with data to be transferred. The buffer preparation routine then updates the last buffer register 112 as
5 shown in block 240.

Consequently, the DMA controller 100 maintains an internal and updated indication of the last buffer which is available for DMA transfer. When the DMA controller 100 decides whether to continue a DMA transfer sequence, it has current information regarding buffer availability. There may be no long latency as typically
10 occurs if the DMA controller periodically polls the memory. Thus, the efficiency of the total DMA transfer may be enhanced since a longer DMA transfer may be achieved. Furthermore, since the DMA controller 100 need not periodically poll the main memory 150 to determine the status of the buffers 160, the bus traffic from the DMA controller 100 may be reduced. This may advantageously allow increased other uses of the bus.
15 Thus, a system employing the disclosed techniques may achieve more efficient performance than prior art systems.

Figure 3 illustrates another embodiment of a system utilizing DMA controller based buffer tracking. Additionally, Figures 4-6 are flow diagrams illustrating operation of one embodiment of the system in Figure 3. Similarly to the system in Figure 1, the
20 system in Figure 3 includes a processor 345, a memory 350, a DMA controller 300, and a DMA target 315. The DMA target 315 may be an audio coder-decoder (CODEC) such as a CODEC compliant with the AC '97 Specification, Revision 1.01, September 10, 1998, published by Intel corporation of Santa Clara, California. The AC '97 specification was

publicly available at the time of filing of this application on the Internet at <http://www.intel.com/pc-supply/platform/ac97>.

As illustrated, the DMA controller 300 may be integrated into a bus controller 340 which also interfaces with two other buses 342 and 344 (e.g., a peripheral components interconnect bus and a low pin count bus). The system of Figure 3 also includes a
5 memory controller 320 which is coupled to the processor 345 by a bus 344, to the memory 350 by a bus 335, to the bus controller 340 by a bus 330, and to a graphics controller 325 by a bus 326.

The embodiment shown in Figure 3 has a buffer descriptor table 355 in the
10 memory 350 and utilizes index values and a buffer descriptor table base register 302 in the DMA controller 300 to reference buffer descriptors. In turn, each buffer descriptor references one of a set of buffers 360 in the memory 350. An enlarged view of one of the buffer descriptors is shown below the memory 350. The buffer descriptor may include a pointer 356 which points to the memory location of the buffer. The descriptor may also
15 include a length field 358 and a command field 354 with an interrupt on complete (IOC) field 359 and a buffer underrun policy (BUP) field 357.

The IOC field 359 may be used to indicate to the DMA controller 300 whether or not it should signal an interrupt upon completion of each transfer. The BUP field 357 may indicate what the DMA controller should do if there is insufficient data to keep
20 passing on to the DMA target 315. For example, if the DMA target 315 generates an audible audio signal, it may be desirable to continue transmitting the last value to avoid an abrupt and undesirable change in the audio. If the DMA target 315 is transmitting data, it may be preferable to interpolate, zero out, or otherwise effectuate and underrun

policy depending on the type of data being transmitted.

Referring now to the flow diagram of Figure 4, the operation of one embodiment of the system in Figure 3 is shown. In block 400, the base address of the buffer descriptor table 355 is written to the buffer descriptor table base register 302. This operation sets a
5 reference point for the index values which will be stored in registers in the DMA controller 300. In block 405, the first set of buffers are prepared for DMA transfer prior to enabling the DMA controller 300.

After these buffers are prepared for DMA transfer, a current index value (CIV) and prefetch index value (PIV) are stored respectively in a CIV register 304 and a PIV
10 register 306 in the DMA controller 300 as illustrated in block 410. The buffer descriptor table 355 in the memory 350 is implemented as a circular buffer with N entries. Accordingly, the DMA control logic 312 may count modulo N using the index values to implement the circular buffer structure. The index value is added to the value stored in the buffer descriptor table base register 302 to obtain the address in memory of the
15 desired buffer descriptor.

Accordingly, the CIV register 304 points to a buffer descriptor 355a which in turn points to a buffer 360a. The PIV register 306 points to a buffer descriptor 355b (subsequent to 355a) which in turn points to a buffer 360b. Since buffer descriptors are used, the memory locations buffers 360a and 360b need not bear any particular relation to
20 each other. The embodiment illustrated in the flow diagrams counts up from zero to N-1; however, decrementing or other appropriate counting techniques could also be used, with comparisons and offsets also being appropriately reversed where necessary.

In step 415, a last valid index (LVI) value is stored in a LVI register 308 in the
42390.P6344

DMA controller 300. The last valid index value indicates the index (the offset into the buffer descriptor table 355 from the buffer descriptor base address stored in the table base register 302) of the last valid descriptor 355c. The last descriptor 355c points to a buffer 360c that is the last one prepared for a DMA transfer. Once the last valid index value is stored in the LVI register 308, the DMA controller 300 may be enabled.

In block 420, a start bit 314, which may be a bit in a control register, is set in the DMA controller 300. When the DMA controller 300 is first starting, both the buffer descriptor indicated by the current index value and the prefetch index value are fetched as indicated in block 425. During continuous operation, the DMA controller will already have the current buffer descriptor and will only need to prefetch the next buffer descriptor. Notably, prefetching is not required; however, when a multiple-tier data structure such as a table using buffer descriptors is used, prefetching may increase efficiency significantly.

In step 430, the buffer indicated by the current buffer pointer (the CIV) is processed. That is, the already fetched buffer descriptor is examined and the pointer to the buffer in memory is extracted as well as the length. The control logic 312 of the DMA controller 300 performs the appropriate number of read or write cycles based on the length of the buffer found in the buffer descriptor. When the transfer cycles are complete, the control logic 312 may take actions in accordance with the command field 354 of the buffer descriptor.

For example, if the interrupt bit or field (IOC 359) is set or enabled, as tested in block 435, the DMA controller signals an interrupt as which may cause a reclamation routine 365 to be executed as indicated in block 440. If the IOC field 359 is not

set/enabled, the control logic 312 proceeds to compare the current index value to the last valid index. If the current buffer is the last buffer available, as tested in block 445, then the controller halts the DMA transfer as indicated in block 450.

If, however, the current index value is not equal to the last valid index, then there are additional blocks to be transferred and the control logic 312 proceeds to block 455 where the current index value is set to the prefetch index value. Since the descriptor for the prefetch index value has already been prefetched, the DMA controller 300 already has the memory address of the appropriate buffer and can continue data transfer. Additionally, the prefetch index value is adjusted (e.g., in this embodiment incremented) to point to the next buffer descriptor as indicated in block 460, and a prefetch of the descriptor at the new prefetch index value is scheduled as indicated in block 465.

Thereafter, the controller returns to block 430 where the (new) current buffer is processed. This process is repeated until the DMA transfer is halted (block 450) due to the exhaustion of the buffer supply (the current index value equaling the last valid index) or until another event interrupts the DMA transfer.

Figure 5 illustrates one embodiment of the reclamation routine 365. The reclamation routine may be executed in response to an interrupt generated at the end of a buffer (IOC bit set), or the reclamation routine may be implemented as a separate process which repeats periodically or which runs continuously. The reclamation routine maintains a head pointer (HEAD) to the last buffer that was freed. The reclamation routine or other software initially establishes the head pointer after at least one buffer is designated for transfer.

In block 500, the current index value is read from the DMA controller 300. If the

current index value is not greater than the head pointer, as tested in block 505, then there are no buffers which have been processed and have not been reclaimed (marked as free for re-use). Therefore, the reclamation routine 365 is finished reclaiming buffers as indicated in block 510. In other embodiments that use a different counting or buffer tracking technique, different comparisons may be used throughout to determine the relationship between the head pointer and the current index value.

If the current index value is less than the head pointer, the reclamation routine 365 marks the buffer pointed to by the head pointer as free as indicated in block 515. Thus, the buffer is released relatively quickly for use by other processes or for reuse in the DMA transfer since the entire DMA transfer need not be completed before the buffer is released. As indicated in block 520, the head pointer is adjusted to point to a next potentially reclaimable buffer (e.g., in this embodiment incremented) and the process returns to reading the current index value from the DMA controller 300.

Since the reclamation process does not happen instantaneously, the current index value may have advanced, allowing additional buffers to be reclaimed. Additional interrupts which may have occurred due to the completion of the processing of other buffers may be ignored since this embodiment of the reclamation routine continues looping and re-loading the current index value until all processed buffers have been reclaimed. This aspect may be particularly advantageous where the DMA controller uses non-reentrant interrupts that may have otherwise been lost to signal the end of each buffer.

Figure 6 illustrates one embodiment of a preparation routine 370. The preparation routine 370 alters the current index value to extend the DMA transfer and include

additional buffers that became available after the DMA transfer was initiated. As indicated in block 600, the tail pointer (TAIL) is set to the current index value. If the next buffer (TAIL + 1) is not free, as tested in block 605, the routine exits in block 610. Either a lack of available buffers, data, or buffer descriptors may prevent the preparation of additional buffers. If those resources needed are available, the tail pointer is incremented in block 615.

Next, a buffer and buffer descriptor are prepared as indicated in block 620. A free buffer may be selected if the DMA transfer is scheduled to read data into memory. A buffer may be filled with data if the DMA transfer is scheduled to transfer data from memory. The memory address of the buffer may be written to the pointer field 356 of the buffer descriptor and the length and special commands may be indicated in appropriate fields as well. Notably, this or other known or otherwise available buffer preparation techniques appropriate for DMA transfers may be used.

After the buffer and buffer descriptor are prepared, the incremented tail pointer is stored in the last valid index register 308. As a result, as soon as the buffer is prepared, the DMA controller 300 may receive current information regarding the last available buffer. In some embodiments, no extraneous bus traffic is required because the software routine updates a value stored in the DMA controller 300 when data for an additional buffer is available for DMA transfer.

Thus, a race free data transfer algorithm using hardware based polling is disclosed. While certain exemplary embodiments have been described and shown in the accompanying drawings, it is to be understood that such embodiments are merely illustrative of and not restrictive on the broad invention, and that this invention not be

42390.P6344 -17-

(

limited to the specific constructions and arrangements shown and described, since various other modifications may occur to those ordinarily skilled in the art upon studying this disclosure.

What is claimed is:

1. A method comprising:

transferring information between a target device and a first buffer which is one of

5 a plurality of buffers, the first buffer being pointed to by a current buffer value stored in a controller;

adjusting the current buffer value stored in the controller to point to a next buffer if the current buffer value is different than a last buffer value stored in the controller; and

10 servicing one of the plurality of buffers utilizing one of the current buffer value and the last buffer value from the controller.

2. The method of claim 1 wherein servicing comprises:

retrieving the current buffer value from the controller;

15 comparing the current buffer value to a head pointer of a buffer list;

marking a buffer pointed to by the head pointer as being processed if the head pointer has not reached the current buffer value; and

adjusting the head pointer to point to a next potentially reclaimable buffer if the head pointer has not reached the current buffer value.

20

3. The method of claim 1 wherein servicing comprises:

preparing a buffer pointed to by an incremented tail pointer; and

storing the incremented tail pointer as the last buffer value in the controller.

42390.P6344

-19-

4. The method of claim 1 wherein the current buffer value is a first index into a buffer descriptor table and wherein the current buffer value points to the first buffer by pointing to a buffer descriptor offset by the first index from a buffer descriptor table base location.
- 5
5. The method of claim 4 further comprising:
- prefetching a next buffer descriptor from the buffer descriptor table.
- 10
6. The method of claim 1 further comprising:
- testing an interrupt field to determine whether the interrupt field is set to a first value to interrupt upon completion of processing the first buffer; and
- generating an interrupt if the interrupt field is set to the first value.
- 15
7. The method of claim 1 further comprising:
- executing a buffer underrun routine according to a command field associated with the first buffer if there are no further buffers available for transfer.
8. A bus agent comprising:
- 20
- a current buffer register for storing a first value indicating a first memory location for a current buffer;
- a last buffer register for storing a second value indicating a second memory location for a last buffer ready for processing; and

control logic coupled to transfer data to or from the current buffer and to update the current buffer register to point to a next buffer unless the first value from the current buffer register is equivalent to the second value from the last buffer register.

5

9. The bus agent of claim 8 wherein the current buffer register contains a first index value and wherein the first index value in the current buffer register indicates the first memory location by pointing to a first buffer descriptor in a buffer descriptor table, the first buffer descriptor being the first index value locations from a buffer descriptor table base.

10

10. The bus agent of claim 9 wherein the last buffer register contains a second index value and wherein the second index value in the last buffer register indicates the second memory location by pointing to a second buffer descriptor in the buffer descriptor table, the second buffer descriptor being the second index value locations from the buffer descriptor table base.

15

11. The bus agent of claim 10 further comprising

a prefetch buffer register, wherein the control logic is coupled to set the current

20

buffer register equal to a value in the prefetch buffer register, to increment the value in the prefetch buffer register, and to schedule a prefetch of a buffer descriptor pointed to by the prefetch buffer register.

12. The bus agent of claim 8 wherein the current buffer register contains a first pointer to a linked list of buffers in memory and wherein the last buffer register contains a second pointer to a last buffer in the linked list of buffers in memory.

5 13. A system comprising:

a bus controller comprising:

a current buffer register for storing a first value indicating a memory location of a current buffer;

10 a last buffer register for storing a second value indicating a memory location of a last buffer, the last buffer being the last buffer which is prepared for processing; and

control logic coupled to transfer data to or from the current buffer and to update the current buffer register to point to a next buffer unless the first value is equivalent to the second value;

15 a processor; and

a memory coupled to the processor, the memory containing:

a plurality of buffers including the current buffer and the last buffer;

at least one software routine which, if executed, causes the system to service

20 at least one of the plurality of buffers utilizing information stored in one of the current buffer register and the last buffer register.

14. The system of claim 13 wherein the memory further contains:

a buffer descriptor table having a plurality of buffer descriptors pointing to the

plurality of buffers, and

wherein the current buffer register contains a first index value, the first index value in the current buffer register indicating the current buffer by pointing to a first buffer descriptor in the buffer descriptor table, the first buffer descriptor being the first index value locations from a buffer descriptor table base.

5

10

15. The system of claim 14 wherein the last buffer register contains a second index value, the second index value in the last buffer register indicating the last buffer by pointing to a second buffer descriptor in the buffer descriptor table, the second buffer descriptor being the second index value locations from the buffer descriptor table base.

15

16. The system of claim 13 wherein the current buffer register contains a first pointer to a linked list of buffers in memory and wherein the last buffer register contains a second pointer to a last buffer in the linked list of buffers in memory.

20

17. The system of claim 13 wherein the bus controller further comprises a prefetch buffer register and wherein the control logic is further coupled to set the current buffer register equal to a value stored in the prefetch buffer register, to increment the prefetch buffer register, and to schedule a prefetch of a prefetch buffer descriptor pointed to by the prefetch buffer register.

18. The system of claim 13 wherein the at least one software routine comprises:

a preparation routine which, if executed, causes the system to perform:

preparing a buffer pointed to by an incremented tail pointer; and

storing the incremented tail pointer in last buffer register.

5

19. The system of claim 13 wherein the at least one software routine comprises:

a reclamation routine which, if executed, causes the system to perform:

retrieving a current buffer value from the current buffer register;

comparing the current buffer value to a head pointer of a buffer list;

10 marking a buffer pointed to by the head pointer as being processed if the head

pointer has not reached the current buffer value; and

adjusting the head pointer to point to a next potentially reclaimable buffer if

the head pointer has not reached the current buffer value.

15 20. A system comprising:

a processor;

a bus agent comprising:

a current index register;

a prefetch index register; and

20 a last index register;

a buffer descriptor base address register;

control logic coupled to transfer data to or from a buffer pointed to by the

current index register, to stop retrieving data if the current index register

42390.P6344

-24-

and the last index register contain equivalent values, to set the current index register equal to the prefetch index register, to increment the prefetch index register, and to schedule a prefetch of a prefetch buffer descriptor pointed to by values stored in the prefetch index register and the
5 buffer descriptor base address register.

a memory coupled to the processor and the bus agent containing:

a table storing a plurality of buffer descriptors

a preparation routine which, if executed, causes the system to perform:

checking an incremented tail index to determine whether the

10 incremented tail index points to a free buffer;

if the incremented tail index points to the free buffer, then

preparing the free buffer by storing data to be retrieved by the

bus master into the free buffer;

storing the incremented tail index in the last index register;

15 a reclamation routine which, if executed, causes the system to perform:

retrieving a current index value from the current index register;

comparing the current index value to a head pointer;

if the head pointer is less than the current index value, then

marking a buffer pointed to by the head pointer as free;

20 incrementing the head pointer;

returning to comparing the current index value to the head
pointer.

21. An article comprising a machine readable medium having stored thereon a plurality of instructions which, if executed by the machine, cause the machine to perform:

transferring information between a direct memory access (DMA) controller and a first buffer which is one of a plurality of buffers, the first buffer being pointed to by a current buffer register in the DMA controller;
adjusting the current buffer register to point to a next buffer if the current buffer register contains a different value than a last buffer register; and
servicing one of the plurality of buffers utilizing information contained in one of the current buffer register and the last buffer register.

10

22. The article of claim 21 wherein the servicing performed by the machine further comprises:

retrieving a first value from the current buffer register;
comparing the first value to a head pointer of a buffer list;
marking the buffer pointed to by the head pointer as being processed if the head pointer has not reached the first value; and
adjusting the head pointer to point to a next potentially reclaimable buffer if the head pointer has not reached the first value.

15

23. The article of claim 21 wherein the servicing performed by the machine further comprises:

preparing a buffer pointed to by an incremented tail pointer; and
storing the incremented tail pointer in last buffer register.

20

42390.P6344

-26-

Abstract

5 A method and apparatus for a race free data transfer algorithm using hardware
based polling. One disclosed method transfers information between a target device and a
buffer which is one of a set of buffers. The buffer is pointed to by a current buffer value
stored in a controller. The current buffer value is adjusted to point to a next buffer if the
current buffer value is different than a last buffer value. One of the set of buffers is
10 serviced utilizing either the current buffer value or the last buffer value from the controller.

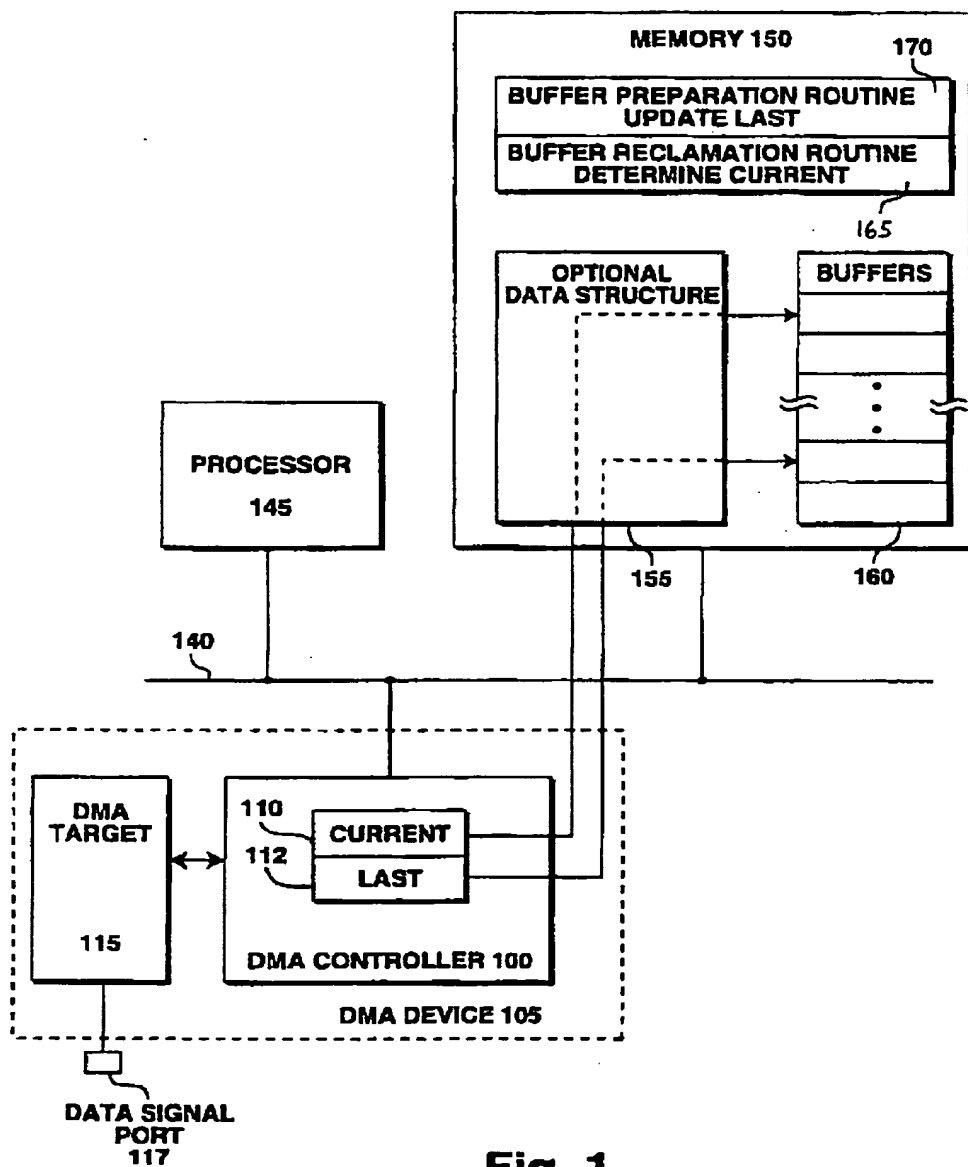
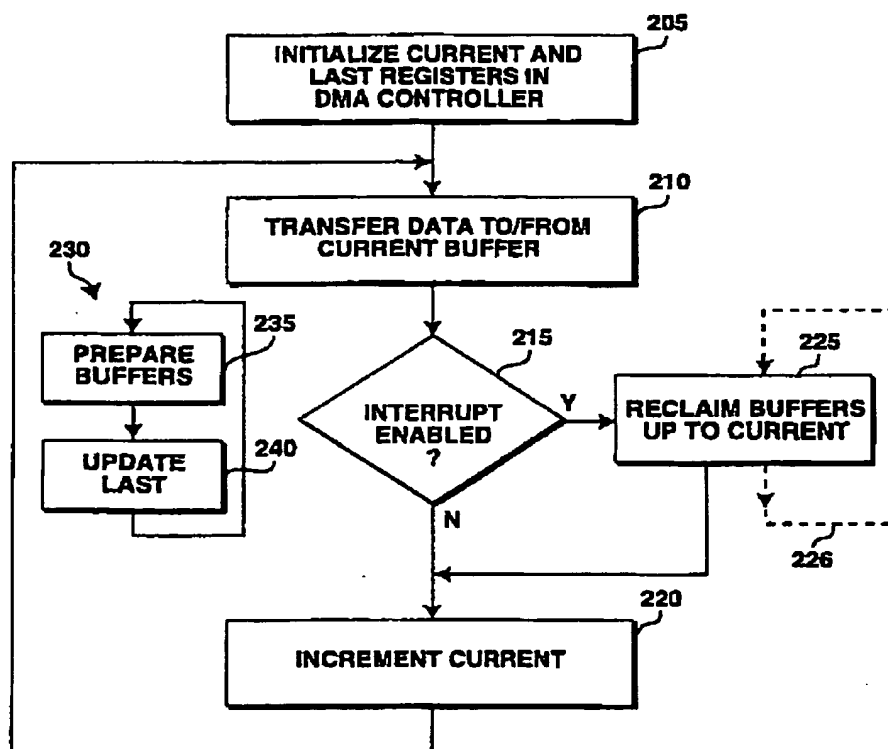


Fig. 1

**Fig. 2**

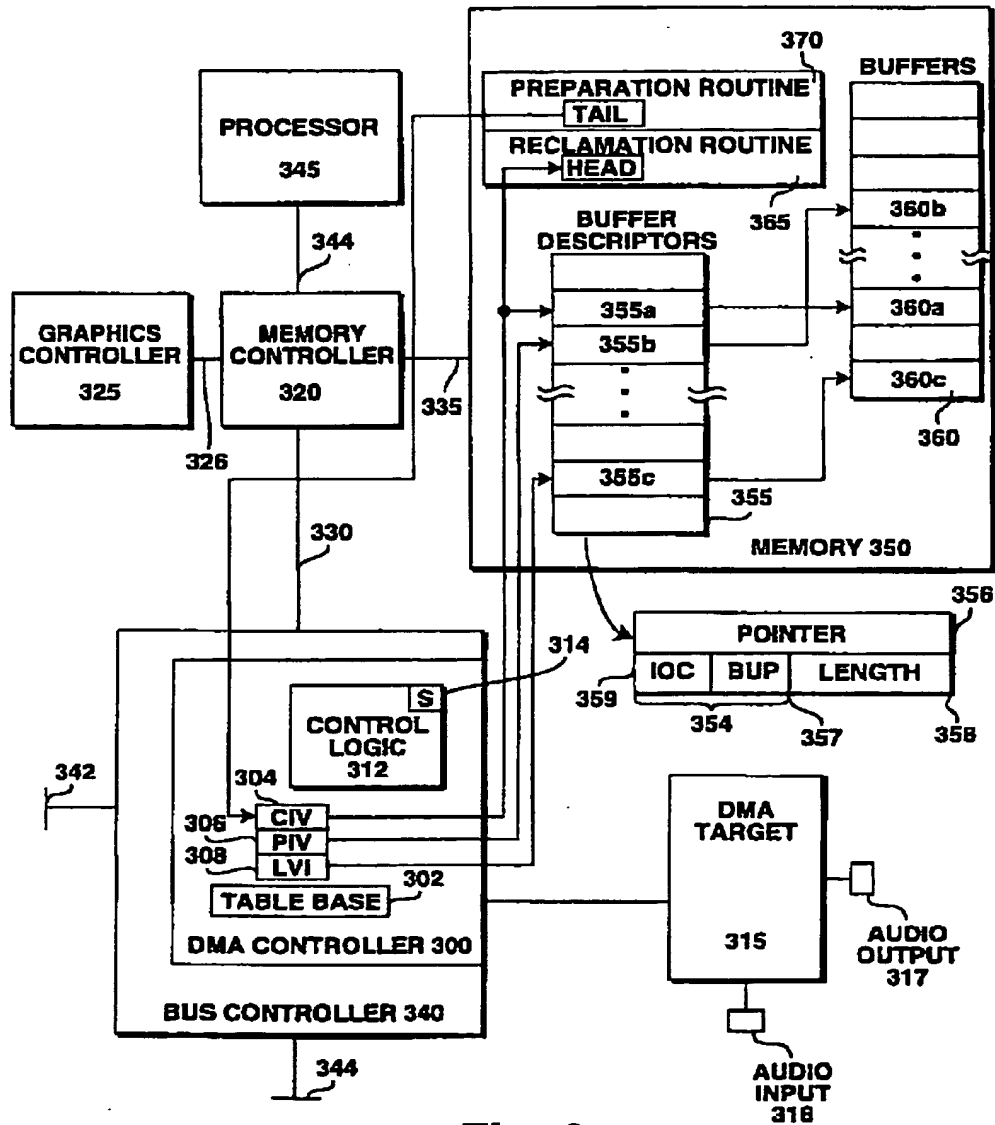
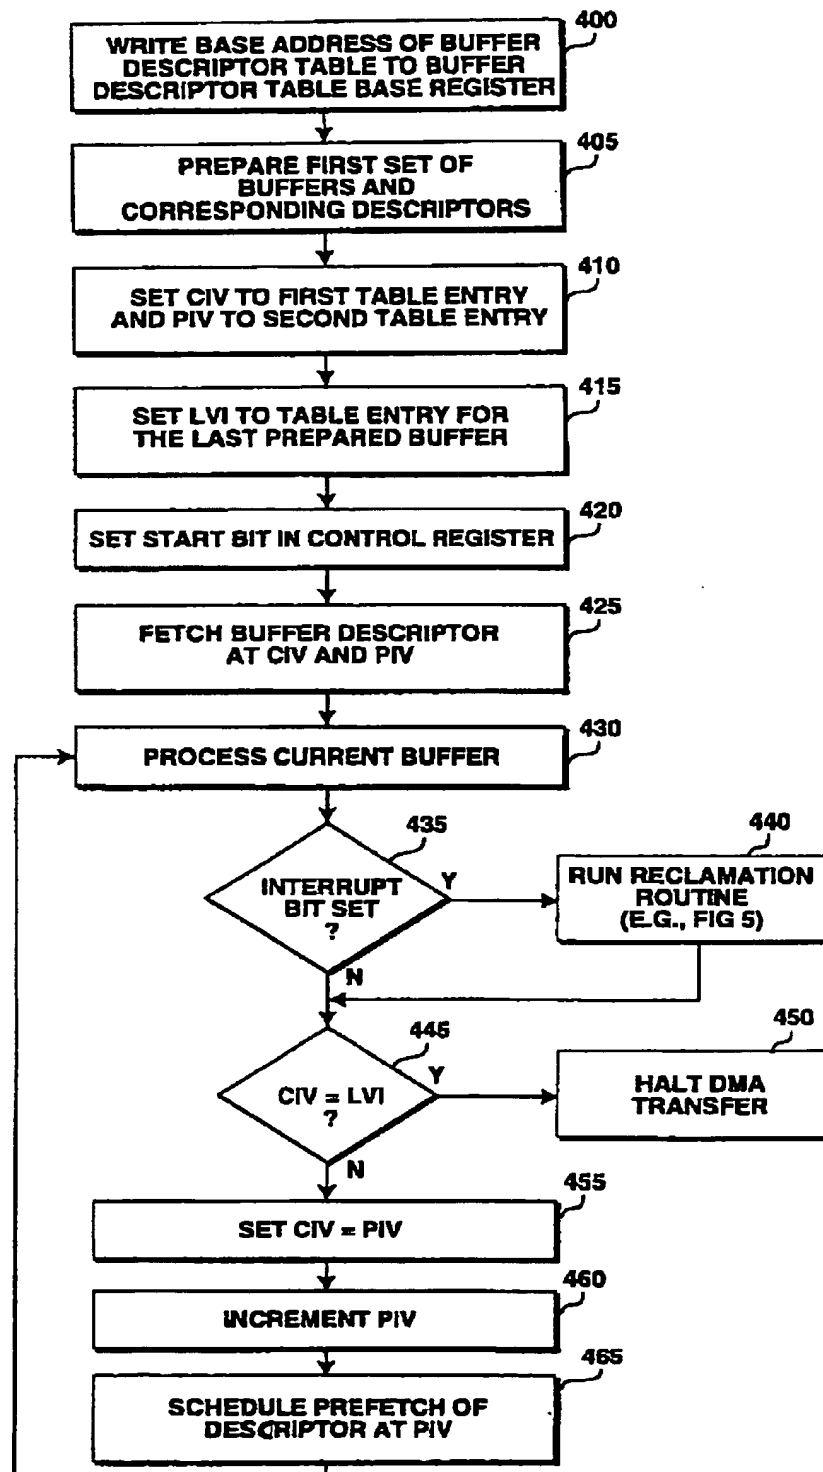


Fig. 3

**Fig. 4**

TOTAL P.05

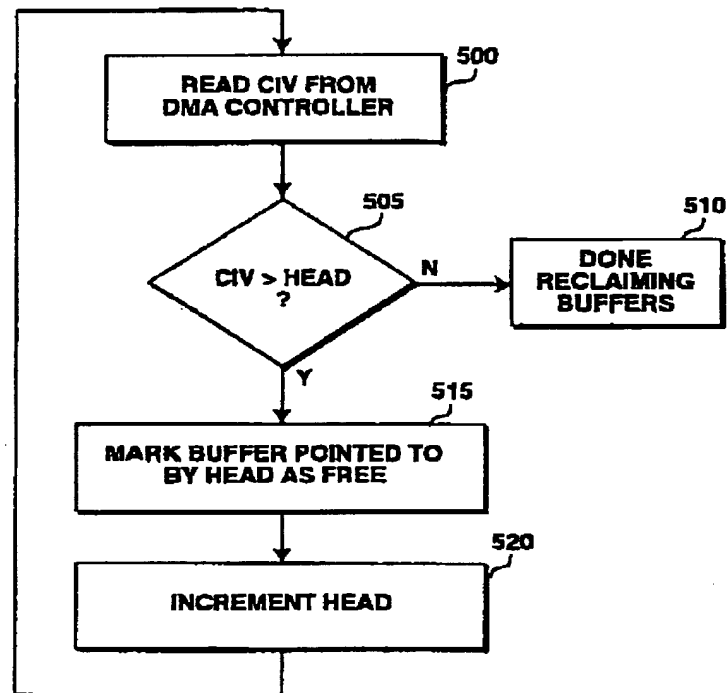



Fig. 5

**Fig. 6**